Good afternoon ladies and gentlemen!

My name is Tom Green and I'm the Program Manager for the Robust Passive Sonar Program, otherwise known as R-P-S.

The focus of R-P-S is to develop innovative, optimal end-to-end processing approaches that exploit tactical acoustic sensors to produce dramatic gains in passive detection against quiet targets operating in shallow water "littoral" environments.

In recent years there has been much emphasis on achieving tactical control of shallow water, littoral regions against quiet diesel electric submarines.

Towards this end, there has been substantial fundamental work aimed at developing advanced processing techniques to exploit the propagation characteristics of these regions.

In addition, new sensors are being developed and deployed that have the necessary characteristics to support many of these processing techniques.

We believe there is an opportunity now to realize significant acoustic gain by applying these and other advanced signal processing techniques to emerging sensors within an end-to-end sonar system context and by so doing to achieve substantial tactical advantage over future submarine threats.

My goal here is to explain why we believe this to be the case, and convey to you where we need your help.

For the benefit of those not familiar with the physics of sound propagation in the ocean, let me provide a quick Acoustics 101 primer.

First of all, sound travels significantly better in water than in air. For example, a car that back-fires in New York might well be heard in Los Angeles.

This should make the problem easy, right?

Not so!

This blessing can also be a curse since at any point in the water many, many such back-firers can be heard and so the problem is not so much one of distinguishing the signal of interest from the noise but instead one of identifying the target of interest from among all the back-firers. In other words, even though you can detect the interesting target signal with your sensor, you just can't distinguish it because it's masked by many uninteresting signals, which we call "clutter".

Second, sound does not travel in a straight line in the ocean, it bends. Unlike my Radar friends who can live in a relatively less complex propagation world, the sonar engineer deals with a medium where the speed of sound and other environmental details vary significantly as a function of position and depth.

This disturbs the acoustic wave front as it propagates through the water and has the effect of bending the acoustic path.

In addition, the sound can bounce off both the ocean surface and the ocean bottom so that a single sound source can be perceived as having several echoes arriving at the listener from a variety of directions.

It turns out that even small changes in the details of the ocean environment, such as temperature and salinity, can significantly impact these complex interactions.

Historically, anti-submarine warfare was concerned with the deep water Soviet, nuclear submarine threat. Detections were made passively, meaning that you listened for sounds generated by a target rather than using active sonar to bounce sound off the target.

As shown in the drawing on the left, the physics of deep-water sound propagation are such that sound may be trapped and travel long distances with little attenuation or bottom interaction. There is also a minimal amount of shipping traffic. Passive acoustics worked well for a long time until threat submarines became quieter.

Furthermore, the downfall of the Soviet Union changed the submarine threat focus. The potential theater of war shifted from the open ocean to shallow water littorals and the world of anti- submarine warfare changed practically overnight.

The current anti-submarine warfare focus is in the coastal environment, where shallow water provides the kind of complicated propagation channel that I described earlier.

Downward refracting sound speed profiles direct the energy towards repeated interaction with, and propagation through, the bottom sediment, which is typically complex and poorly known. The shallow water environment is highly dynamic, and source motion and maneuvering become important.

Because shallow water operations can be close to home, shorter endurance diesel electric submarines can be an ideal platform. As a result, many nations are acquiring or already own these relatively low cost submarines.

In addition, the diesel submarine can be extremely quiet when running on battery and trends are for increasingly quieter signatures.

The coastal environment also means a higher density of shipping traffic, which is the dominant source of low frequency noise.

Measurements show that heavy shipping densities typical of littoral regions can produce ambient noise backgrounds that are 20 dB or more louder than wind-wave noise, which is the dominant noise source at higher frequencies.

What's the impact of such a 20 dB loss in signal-to-noise ratio?

While it depends on the environment, it can make the difference between detecting or not detecting a target at a tactically significant range.

For example, in the Strait of Korea, a 20 dB reduction in sensitivity can reduce 60 km range detection to 5 km range detection.

In other regions and environments, the differences in detection can be equally as dramatic.

Because acoustic energy in the ocean is coherent, that is, phase relationships are preserved over time and space, arrays of acoustic hydrophones can be used to provide improved sensitivity to the signal over a single hydrophone and to enable directional discrimination.

Coherently processing acoustic array data to provide directional sensitivity is known as beam forming.

"Conventional" beam forming assumes that the target signal impinges on the array from a single direction, or relative bearing.

For the deep-water problem, conventional beam forming applied to a linear array of hydrophones is an effective means of processing acoustic array data to localize sound sources in bearing. However, the conventional beam former is not well matched to the complex multi-path environment typical of the littoral, since a single sound source can actually appear to arrive on many conventionally formed beams.

In addition, the conventional beam former is incapable of resolving two sound sources at the same bearing, which as we will see, is essential for detecting quiet targets in the presence of loud surface ships.

The match field processor is better suited for coping with, and actually exploiting, this complicated propagation channel.

Instead of assuming a single incident direction for the source energy and forming beams for different bearings relative to the array, the matched-field processor utilizes information about the multi-path propagation characteristics of the channel to form beams in bearing, range and depth.

In essence, matched-field processing divides up the ocean into many "volume cells" and estimates the amount of acoustic energy originating from each one.

Matched field processing offers another important advantage.

I mentioned earlier that the primary noise source in shallow water is often shipping traffic.

This shipping noise is spatially discrete and located on the ocean surface.

The ability of matched field processing to localize in range, depth and bearing offers the possibility of resolving substantial shipping noise contributors from targets of interest.

This ability to resolve and consequently suppress shipping noise interference is significantly enhanced via space-time adaptive processing to adapt the matched field processing beam former to the noise environment.

That's the good news.

However, there are challenges that remain.

Remember that to carry out matched-field processing we require some knowledge of the environmental details of the acoustic channel, for example the sound speed profile.

Remember also that these details change with time and with location in ways that can be significant.

Thus, an effective matched-field processor must be robust to this environmental variability.

In addition, adaptive matched-field processing relies on estimating the statistics of the noise directly from the acoustic data.

This requires a certain degree of "stationarity" in the scene or a period of time for which things don't change in a statistically significant way.

For standard, what we call "full-rank" adaptive techniques, this required stationarity time grows with the size of the array.

Meeting this stationarity requirement for large arrays in shallow water environments can be particularly challenging because of the dynamic maneuvers that characterize littoral operations.

Ensuring matched-field processor robustness to both environmental uncertainty and source-receiver dynamics is essential to achieving the goals of R-P-S.

Recognizing the importance of matched field processing to effective acoustic processing in littoral environments, DARPA conducted the Santa Barbara Channel Experiment. The objective of the experiment was to understand, in a controlled environment, the fundamental limits on the performance of adaptive matched field processing, and to develop signal-processing methods that overcome the challenges just mentioned.

The experiment deployed five, instrumented vertical line arrays, with 30 hydrophones each, that collected data on a variety of surface and submerged acoustic sources.

A number of very important and promising results came out of the experiment and the subsequent analysis.

For example, we show here Matched Field Processing ambiguity surfaces, which are plots of estimated acoustic intensity versus range and depth, from a single array.

The data set involved a research vessel, called the Acoustic Explorer that towed a submerged sound source at a depth of 30 meters.

The source emitted a continuous sequence of 9 narrowband tones. The left-hand plot shows the result for conventional Matched Field Processing, that is to say it is matched field processing that does not adapt to the noise characteristics in the scene.

This ambiguity surface has a peak at the correct position of the Acoustic Explorer, but the high side lobes from the naturally radiated noise of the Acoustic Explorer obscure the target.

The right-hand plot shows the result for adaptive matched field processing using the same data. In this plot, the peak of the lower energy target, the one that is submerged, is clearly seen in the presence of the surface ship, even though the ship is at the same bearing and range and is separated from the target by only 30 meters in depth.

Similar results were obtained for the naturally radiated signature of a submarine that participated in the experiment when heavy shipping traffic was present.

The standard approach to adaptive processing is to represent the acoustic field by a vector whose dimensionality is equal to the number of array elements. This is called "full-rank" processing. Properly adapting to the noise statistics for a full rank technique requires the estimation of a covariance matrix whose dimensionality is equivalent to the number of elements in the array.

Full rank adaptive matched field processing can produce excellent localization accuracy, but for large arrays the number of observations required to properly estimate the necessary statistics can be prohibitive.

This architecture is also vulnerable to environmental mismatch and losses from target motion. Fortunately, the true dimensionality of the noise field of interest is often much lower than the dimensionality of the array and this reduced dimensionality can be exploited.

In this chart we consider the application of modal decomposition and modal filtering which addresses these issues by reducing dramatically the rank of the covariance matrix involved.

In addition, one can design the modal filter so that energy from the surface is removed, amounting to doing surface ship filtering.

Both plots are adaptive matched field processing results calculated for a single array involving the Acoustic Explorer and the towed source.

The left-hand plot shows the full rank adaptive matched field processing output, which uses adaptivity to exploit the depth separation of the tow ship and the submerged source. This result is possible for the particular vertical array we used, but the issues I just described arise when a larger array is involved.

The right-hand plot shows a reduced-rank processor for the same data.

The submerged source is still present, but the mode filter has removed the tow ship.

In addition, the Adaptive Matched Field Processing beam has broadened. This broadening provides increased robustness to mismatch and motion, and also relaxes the requirement to compute the Matched Field Processing output along very fine range grids.

A long observation time increases signal gain and helps satisfy the snapshot requirement for adaptive processing, but it's hindered by uncompensated scene motion. Motion compensation essentially adjusts the data snapshots so that a target appears stationary. It has the added advantage that it simultaneously de-focuses interference sources that are not moving along the target track.

We achieved motion compensation by utilizing a velocity hypothesis to determine the magnitude and phase adjustment required at each hydrophone in the array, and for each observation time. This is equivalent to propagating the target at each time to a fixed "focus" position.

The plot on the left shows the adaptive matched field-processing surface for a 300 second observation time with no motion compensation applied. The Acoustic Explorer moves 750m during this time period and that results in the smearing and elongation of the target peak.

The right-hand plot shows the adaptive matched field-processing surface with motion compensation applied. The result shows the concentration of source energy at the focus and the simultaneous de- focusing of energy in the side lobes.

Collectively, these results constitute significant progress towards developing robust methods for Adaptive Matched Field Processing in littoral environments.

They are the foundation for the new DARPA Robust Passive Sonar program, or R-P-S.

R-P-S intends to extend the results obtained to tactical sensors, thereby providing a revolutionary new capability for tactical sensors.

The TB-29 is an emerging submarine towed array that is well suited to support many of the advanced signal processing methods required for littoral anti-submarine warfare.

Other sensors available to us are the SURTASS Twin Line array and emerging prototype arrays.

The basic objective of R-P-S is to provide significant improvement in figure-of-merit performance for tactical anti-submarine warfare in the littoral over baseline systems.

I have suggested some possibilities for how to realize these gains based on important results derived from the Santa Barbara Channel Experiment, but we need your ideas, too!

For instance, how can we exploit external sensors to observe surface ship locations, and apply the knowledge gained to the tactical anti-submarine warfare scenario? How can we carry out in situ environmental calibration with sufficient fidelity to enable advanced processing methods? How can we cope efficiently with the high- dimensional search problem that is presented by a high-resolution, tactical sensor? These and other questions must be answered to enable the gains that are possible.

What we are looking for in the R-P-S program is to investigate new methods for target ranging and for suppression of discrete surface ship interference.

Our general approach is to:

Conduct systems analysis and assessments to determine how these new processing techniques may translate into new sensor concepts, new platform tactics, or new system architectures,

Extend what we learned in the Santa Barbara Channel Experiment to emerging and future tactical platforms and tactical systems,

Integrate adaptive processing algorithms and/or other novel processing techniques into an end-to-end prototype architecture that optimizes processing gains,

And go to sea and test these algorithms in the real world!

We need your help! I believe the Robust Passive Sonar program can produce dramatic gains in shallow water passive acoustic signal processing, but your knowledge and ideas are essential.

In the coming months, we will publish a call for proposals.

Let your imaginations get the best of you.

With DARPA, the only limits are those that Mother Nature places on us and even they are negotiable!

I will be here the remainder of the conference and look forward to meeting as many of you as possible.

Thank you very much for your attention.